

Precise Augmented Reality Enabled by Carrier-Phase Differential GPS

Daniel P. Shepard, Kenneth M. Pesyna, Jr., and Todd E. Humphreys
The University of Texas at Austin

BIOGRAPHIES

Daniel P. Shepard is pursuing a Ph.D. in the Department of Aerospace Engineering and Engineering Mechanics at The University of Texas at Austin, where he also received his B.S. in 2011. He is a member of the UT Radionavigation Lab. His research interests are in GNSS security, estimation and filtering, and guidance, navigation, and control.

Kenneth M. Pesyna, Jr. is pursuing a Ph.D. in the Department of Electrical and Computer Engineering at The University of Texas at Austin where he received his M.S. in 2011. He received his B.S. in Electrical and Computer Engineering from Purdue University in 2009. He is a member of the UT Radionavigation Laboratory and the Wireless Networking and Communications Group. His research interests lie in the areas of indoor navigation and precise low-power mobile positioning.

Todd E. Humphreys is an assistant professor in the department of Aerospace Engineering and Engineering Mechanics at The University of Texas at Austin, and Director of the UT Radionavigation Laboratory. He received a B.S. and M.S. in Electrical and Computer Engineering from Utah State University and a Ph.D. in Aerospace Engineering from Cornell University. He specializes in applying optimal estimation and signal processing techniques to problems in radionavigation. His recent focus is on radionavigation robustness and security.

ABSTRACT

A prototype precise augmented reality (PAR) system that uses carrier phase differential GPS (CDGPS) and an inertial measurement unit (IMU) to obtain sub-centimeter level accurate positioning and degree level accurate attitude is presented. Several current augmented reality systems and applications are discussed and distinguished from a PAR system. The distinction centers around the PAR system's highly

accurate position estimate, which enables tight registration, or alignment of the virtual renderings and the real world. Results from static and dynamic tests of the PAR system are given. These tests demonstrate the positioning and orientation accuracy obtained by the system and how this accuracy translates to remarkably low registration errors, even at short distances from the virtual objects. A list of areas for improvement necessary to create a fully capable PAR system is presented.

I. INTRODUCTION

The concept of virtual reality has been around since shortly after the advent of the computer. People have been fascinated with the idea of designing a world of entirely their own creation that can be visualized, explored, and interacted with. Initially, attempts at creating such environments were greatly limited by the technology of the time. Creating a realistic-looking, real-time representation of a virtual environment that a user can interact with requires fast processors, high screen resolution, a significant amount of memory, and large communication bandwidths. Technological advances in the last couple of decades have enabled many applications of virtual reality including immersive training simulations and gaming.

In recent years, the concept of augmented reality has garnered significant attention. While virtual reality seeks to replace the real world with a simulated one, augmented reality blends together real world and virtual elements typically through context relevant visuals. One can image a continuum of perception, with the real world on one end and virtual reality on the other. Augmented reality seeks to place the user's perception somewhere in the middle of the continuum with the goal of complementing the real world with virtual elements.

Augmented reality applications often require knowledge of the platform's position and orientation display context sensitive information. When positioning

in a global coordinate system is required, all current consumer platforms running augmented reality applications are only capable of obtaining the user’s location to within a few meters. This level of positioning uncertainty severely limits potential applications because accurate alignment of the real and virtual worlds cannot be obtained at close distances.

This paper presents a prototype augmented reality system that is capable of obtaining sub-centimeter level accurate positioning and degree level accurate attitude. The paper starts with an overview of current augmented reality applications and the techniques used to blend the real and virtual worlds together. The concept for the prototype system is then presented. Next, a description of the prototype system and navigation filter are given. Finally, results from tests of the prototype system are presented and remaining technical challenges are discussed.

II. AUGMENTED REALITY OVERVIEW

Augmented reality has a presence in many fields including medicine, sports, art, architecture, tourism, entertainment, and marketing. In sports, the first down line on a live TV broadcast of a football game is an example of augmented reality. This technology uses a combination of visual cues from the field as well as the location of the cameras producing the video feed [1]. In marketing, augmented reality is used to help consumers make shopping decisions, such as the augmented reality system employed in Lego kiosks where the Lego product is displayed fully assembled on top of the box when held up in front of a smart-phone camera [2]. This application uses visual tags on the box to orient the product correctly. Augmented reality can also be used in tourism, where augmented reality applications, as an example, can be used to translate signs from one language to another. Fig. 1 shows the smart-phone application “Word Lens” translating a sign from Spanish to English.

Augmented reality applications use techniques that can be classified into two primary, yet overlapping, categories: image processing and so-called pose or position and orientation. Image processing techniques are those as described in the examples above, which use information from the video stream to properly orient and position computer-rendered virtual objects into the video feed. In many cases, fiduciary mark-



Fig. 1. A smart-phone running the Word Lens application that is translating a sign from Spanish to English [2].

ers, referred to as cues, are placed in the field of view of the camera to allow for proper alignment of the virtual and real objects. These markers are typically of known size and shape. By identifying the observed marker of known size and shape, the AR system can position and orient the camera relative to the marker. In other cases, color palettes are used to determine colors over which virtual objects can be placed. The first-down line in football, which must be placed on the field and avoid being overlaid on referees or players, is a good example of this technique. Green screens are another.

Pose techniques are those that employ the use of position and orientation to properly register the virtual information onto the real world. This technique uses sensor information to obtain position and orientation and properly align or “register” virtual objects with the real world. In some setups, orientation and position of the user is done using visual cues, as with fiduciary markers. However, this method requires some setup: the visual cues must be pre-placed and positioned in the real environment. Another technique is to use the absolute user position and orientation rather than the relative position and orientation provided by visual cues. Systems employing absolute positioning techniques require the absolute position and orientation of the user, the virtual objects, and possibly the real objects to properly create the augmented reality environment. The advantage of this approach is that visual markers are no longer needed.

Absolute positioning techniques require a shared coordinate system between the user and the objects in the environment. This shared coordinate system is often the Earth-centered, Earth-fixed (ECEF) coordinate system. Augmented reality applications em-



Fig. 2. The Star Walk application is being used to highlight constellations in the night sky [3].

ploying absolute positioning most often use GPS for positioning and an inertial measurement unit (IMU) for orientation. Astronomy is one application where absolute-positioning-based augmented reality is used. Applications like “Star Walk,” shown in Fig. 2, created by Vito Technology allow users to point their smart-phone or tablet at the sky and highlight constellations. The application uses coarse position and orientation information provided by the on-board sensors to calculate the section of the sky at which the device is pointed [3].

In the case of the “Star Walk” application, the real-world objects, the stars, are far from the user. Consequently, the 3-to-10 meter positioning accuracy provided by pseudorange-based GPS is adequate for this application. However, there are many applications where this positioning accuracy is far from adequate. In the construction industry, ground excavation is often performed when, for example, laying the foundation for a new building or building a new road. Before excavation begins, workers must determine the locations of pre-existing subsurface pipelines and power cables. Field operatives must determine the location of these underground utilities from maps of the underground infrastructure and mark their location to ensure excavation does not damage any pipes or cables. This process can be a cumbersome task. If an augmented reality system capable of determining a user’s location and attitude with a high degree of accuracy is available, then this process can be simplified by visually overlaying the subsurface infrastructure maps onto a live video feed for the field operative. Figure 3 demonstrates such a system developed by a group of



Fig. 3. An image produced by an augmented reality system developed by researchers at the University of Nottingham. This image shows subsurface utilities visualized in their correct location relative to the real world image in the background [4].

researchers from the University of Nottingham [4].

In applications such as infrastructure overlay, meter-level errors in user positioning are unacceptable because the user needs to determine the exact location of the virtual objects by walking up to the objects. There exists, in these situations, a need for so-called “precise” augmented reality. In these systems, user positioning and orientation is accurate at the sub-centimeter and degree levels respectively. This positioning accuracy can be obtained by what is known as carrier-phase differential GPS, and the orientation accuracy can be obtained with a medium-grade IMU.

III. PRECISE AUGMENTED REALITY (PAR)

Precise augmented reality (PAR) can be defined as an augmented reality system exploiting highly accurate position and orientation information. This paper presents a PAR system that makes use of sub-centimeter level accurate positioning, enabled by carrier-phase differential GPS (CDGPS), and degree level orientation, provided by a medium-grade IMU. As discussed briefly in Sec. II, existing work in this area consists of a system built for subsurface data visualization [4]. This system used a CDGPS receiver built by Leica Geosystems to provide the position estimate. Additionally, a group from Columbia University built a head-mounted augmented reality system, shown in Fig. 4, for the general purpose of



Fig. 4. An augmented reality system developed by researchers at Columbia University for providing users with information about their surroundings [5].

providing users information about their surroundings as they navigate an unfamiliar urban environment [5]. This system, however, used a weaker form of correction-based precise positioning called differential GPS, which produced meter-level positioning errors.

Neither of the previously mentioned systems take advantage of the gains that can be obtained by coupling the IMU and the GPS measurements together. The system presented in this paper goes beyond this prior work by integrating the CDGPS and IMU measurements within an extended Kalman filter. The resulting accurate pose estimated permits tight registration of virtual objects.

IV. PROTOTYPE PAR SYSTEM

A prototype PAR system was designed and built to demonstrate the potential of precise augmented reality. A picture of the inside of the system with labels for the various components is shown in Fig. 5. The system is comprised of a GPS receiver, an IMU, a webcam, a single-board computer (SBC), a GPS antenna, and a lithium-ion battery. Detailed descriptions of the GPS receiver, IMU, and webcam used are given in Sec. IV-A, IV-B, and IV-C respectively. This sensor package is strapped to the back of a tablet PC, which collects the data from the GPS receiver and IMU and the webcam footage for post-processing. In

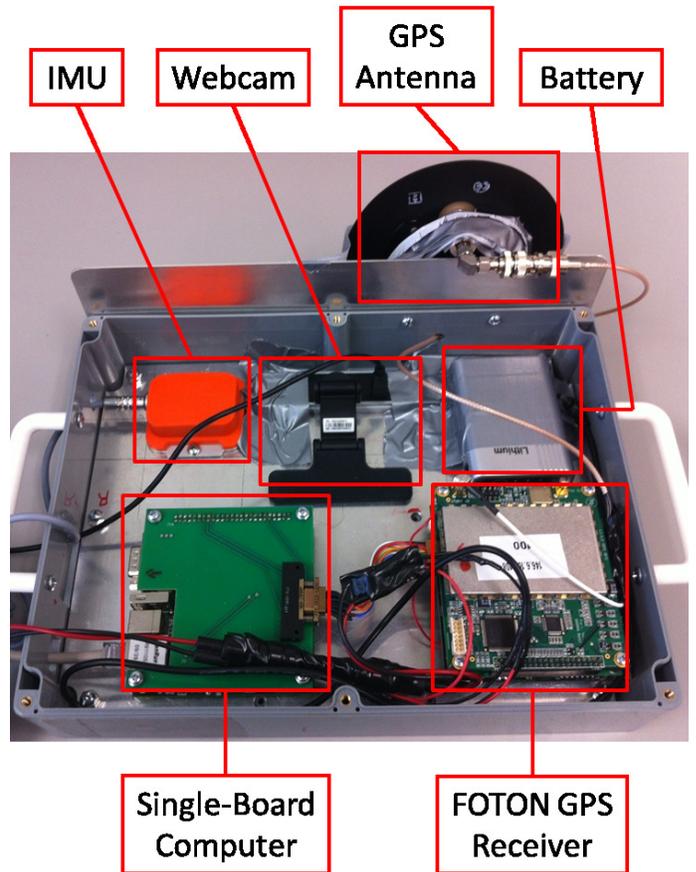


Fig. 5. The inside of the prototype PAR system with labeled components.

a real-time version of the system, the tablet PC would take the form of a “window” into the AR environment; a user looking “through” the tablet would see an augmented representation of the real world on the other side of the tablet. A picture of the fully assembled system is shown in Fig. 6.

A. GPS Receiver

The GPS receiver used for the prototype PAR system was the FOTON software-defined GPS receiver. The work-horse of the receiver is a digital signal processor (DSP) running the GRID software receiver developed by The University of Texas at Austin and Cornell University. This software receiver was originally developed for Ionospheric monitoring. As such, it has a scintillation robust PLL and databit prediction capability, which both help to prevent cycle slips [6]. The receiver is also dual-frequency and currently capable of tracking L1 C/A and L2C, but only the L1 C/A signals were used in the prototype PAR system. Observables and navigation solutions can be output at a



Fig. 6. The assembled prototype PAR system.

configurable rate which was set to 5 Hz. The receiver communicates with the tablet over Ethernet via the SBC.

B. IMU

The IMU used for the prototype PAR system was the MTi developed by Xsens. This IMU is a complete gyro-enhanced attitude and heading reference system (AHRS). It houses four sensors, (1) a 3D magnetometer, (2) a 3D gyro, (3) a 3D accelerometer, and (4) a thermometer. The MTi also has a DSP running a Kalman filter (KF) that determines the attitude of the MTi relative to the north-west-up (NWU) coordinate system. This KF determines attitude by ingesting temperature-calibrated (via the thermometer and high-fidelity temperature models) magnetometer, gyro, and accelerometer measurements to determine magnetic north and the gravity vector, which is sufficient for full attitude determination. This estimate of orientation is accurate to better than 2° RMS, during dynamic operation.

In addition to providing orientation, the MTi also provides access to the highly-stable, temperature-calibrated measurements. The configurable output rate of the MTi was set to 100 Hz. In order to obtain a time stamp for the MTi data, the MTi measurements were triggered by the FOTON receiver which also reported the GPS time the triggering pulse was sent.

C. Webcam

The webcam used for the prototype PAR system was the FV Touchcam N1. The Touchcam N1 is an HD webcam capable of outputting 720P at 22 fps or WVGA at 30 fps. The Touchcam N1 also has a wide angle lens with a 78.1° horizontal field of view.

V. PAR EKF

The goal of the navigation filter for this PAR system was to obtain sub-centimeter level accurate positioning and degree level accurate attitude of the webcam. The IMU chosen for the prototype PAR system already provided attitude to the desired accuracy through its own Kalman filter. Therefore, the filter designed here only needed to satisfy the positioning requirement. For a first prototype, it was assumed that the attitude and angular velocity provided by the IMU were perfect for purposes of designing the filter. Although this assumption is inaccurate, the effect of errors in attitude and angular velocity on the estimate of position is negligibly small for this system.

The navigation filter designed for this PAR system was a tightly-coupled INS/GPS enhanced Kalman filter (EKF) implemented as a square-root information filter. A block diagram of this filter is shown in Fig. 7. The filter can be divided into three primary components: (1) CDGPS measurement update, (2) inertial navigation system (INS) propagation step, and (3) augmented reality overlay. These primary components are described in detail in the sections that follow.

The state for the EKF is

$$\vec{X} = \begin{bmatrix} \vec{x} \\ \vec{N} \end{bmatrix} = \begin{bmatrix} \vec{r} \\ \dot{\vec{r}} \\ \vec{b} \\ \vec{N} \end{bmatrix} \quad (1)$$

where \vec{X} is the state vector, \vec{x} is the real-valued part of the state vector, \vec{N} is the carrier phase integer ambiguity vector, \vec{r} and $\dot{\vec{r}}$ are the ECEF position and velocity vectors of the IMU respectively, and \vec{b} is the accelerometer bias vector. This state differs from that of the GPS antenna and the webcam lens in position and velocity due to the distance between the system components and the rigid-body rotation of the entire

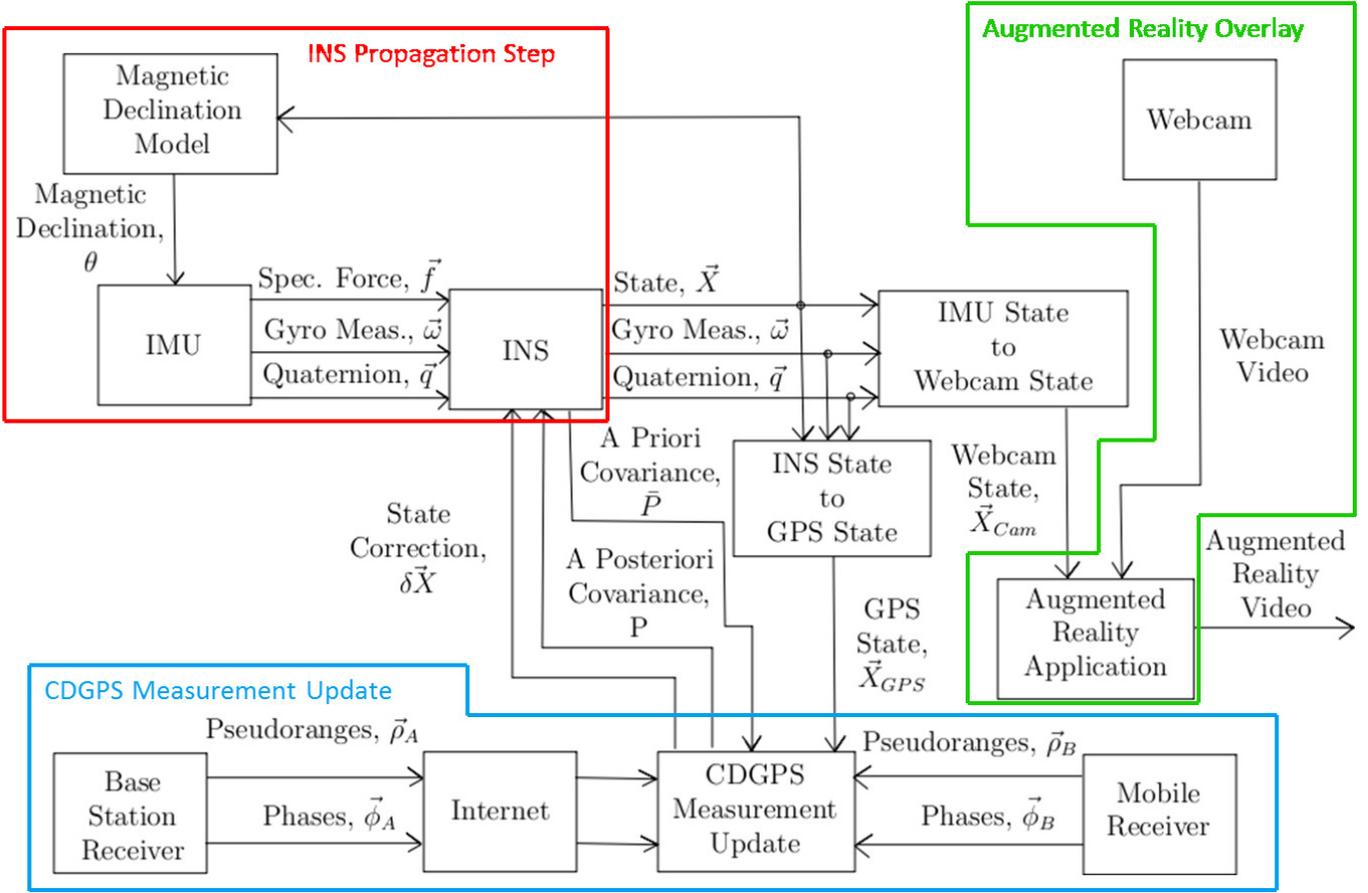


Fig. 7. A block diagram of the PAR EKF. The blocks corresponding to the INS propagation step (red), CDGPS measurement update (blue), and augmented reality overlay (green) are highlighted in the figure.

system. This is accounted for through the following transformation for the GPS antenna state:

$$\begin{aligned}\vec{r}_{GPS} &= \vec{r} + R_{ECEF}^{IMU} \vec{r}_{GPS/IMU} \\ \dot{\vec{r}}_{GPS} &= \dot{\vec{r}} + R_{ECEF}^{IMU} (\vec{\omega} \times \vec{r}_{GPS/IMU})\end{aligned}\quad (2)$$

where \vec{r}_{GPS} and $\dot{\vec{r}}_{GPS}$ are the ECEF position and velocity vectors of the GPS antenna respectively, R_{ECEF}^{IMU} is the rotation matrix from the IMU frame to the ECEF frame that is derived from the quaternion output by the IMU and the current best position estimate of the IMU, and $\vec{r}_{GPS/IMU}$ is the position vector of the GPS antenna relative to the IMU in the IMU frame. To obtain the webcam lens position and velocity, replace \vec{r}_{GPS} , $\dot{\vec{r}}_{GPS}$, and $\vec{r}_{GPS/IMU}$ in Eq. 2 with \vec{r}_{CAM} , $\dot{\vec{r}}_{CAM}$, and $\vec{r}_{CAM/IMU}$. In order for the EKF to be accurate to the sub-centimeter level, the relative position vectors of the components must be known to millimeter-level accuracy.

A. CDGPS Measurement Update

The CDGPS measurement update is responsible for determining the accurate position of the system and correcting for IMU drift. The CDGPS algorithm requires observable measurements (carrier phase and pseudorange) from two separate GPS receivers. The first receiver is a reference or base-station receiver that is located at a pre-surveyed position. The second receiver is the mobile receiver, which is located on the PAR system. The location of the base station receiver's antenna must be known with high accuracy, since the CDGPS algorithm only provides an accurate position of the mobile receiver's antenna relative to the base station receiver's antenna.

A.1 Measurement Model

The CDGPS algorithm uses single frequency (L1 C/A) double-differenced measurements of carrier phase and pseudorange to determine position to sub-

centimeter level accuracy. These double-differenced measurements are formed as follows:

single-differenced measurements: The pseudorange, $\rho^i(k)$, and carrier phase, $\phi^i(k)$, measurements from each receiver taken at the same instant in time are matched up by satellite and differenced between the two receivers to obtain single-differenced pseudorange and carrier phase measurements. These single-differenced pseudorange and carrier phase measurements are denoted as

$$\begin{aligned}\Delta\rho_{AB}^i(k) &= \rho_A^i(k) - \rho_B^i(k) \\ \Delta\phi_{AB}^i(k) &= \phi_A^i(k) - \phi_B^i(k)\end{aligned}\quad (3)$$

The subscripts A and B represent the base station and mobile receivers respectively. Assuming the base station and mobile receivers are relatively close to one another (within about 10 km) and the measurements were taken at roughly the same time, these single-differenced measurements have removed the errors in the measurements that are common to the satellite. These errors are: (1) Ionospheric and Tropospheric delays, (2) satellite clock errors, and (3) the initial broadcast carrier phase of the signal.

double-differenced measurements: Of the satellites tracked by both receivers, one satellite is chosen as the ‘‘reference’’ satellite (satellite 0). The single-differenced pseudorange and carrier phase measurements corresponding to the reference satellite are subtracted from all the other single-differenced pseudorange and carrier phase measurements to create double-differenced pseudorange and carrier phase measurements. These double-differenced pseudorange and carrier phase measurements are denoted as

$$\begin{aligned}\nabla\Delta\rho_{AB}^{i0}(k) &= \Delta\rho_{AB}^i(k) - \Delta\rho_{AB}^0(k) \\ \nabla\Delta\phi_{AB}^{i0}(k) &= \Delta\phi_{AB}^i(k) - \Delta\phi_{AB}^0(k)\end{aligned}\quad (4)$$

These double-differenced measurements have removed the effect of receiver clock bias and, if the receiver is designed properly, the ambiguities on the carrier phase measurements have now become integers. These integer ambiguities can be determined much faster than their real-valued counterparts through the use of the integer constraint.

A detailed derivation of the double-differenced measurement model for the carrier-phase can be found in Ref. [7]. The double-differenced pseudorange model is derived in the same way with the ambiguity term removed. The double-differenced measurement models are given as

$$\begin{aligned}\nabla\Delta\rho_{AB}^{i0}(k) &= \nabla\Delta r_{AB}^{i0}(k) + \nabla\Delta\nu_\rho \\ \lambda_{L1}\nabla\Delta\phi_{AB}^{i0}(k) &= \lambda_{L1}N_{AB}^{i0} + \nabla\Delta r_{AB}^{i0}(k) + \nabla\Delta\nu_\phi\end{aligned}\quad (5)$$

where $\nabla\Delta\nu_\rho$ and $\nabla\Delta\nu_\phi$ are the white Gaussian measurement noises, λ_{L1} is the wavelength of the GPS L1 center frequency, N_{AB}^{i0} is the integer ambiguity, and

$$\begin{aligned}\nabla\Delta r_{AB}^{i0}(k) &= (\|\vec{r}_{SV}^i(k) - \vec{r}_A\| - \|\vec{r}_{SV}^i(k) - \vec{r}_B(k)\|) \\ &\quad - (\|\vec{r}_{SV}^0(k) - \vec{r}_A\| - \|\vec{r}_{SV}^0(k) - \vec{r}_B(k)\|)\end{aligned}\quad (6)$$

where $\vec{r}_{SV}^i(k)$ is the ECEF position of satellite i , \vec{r}_A is the location of the base station receiver, and $\vec{r}_B(k) = \vec{r}_{GPS}(k)$.

The measurement model in Eq. 5 is linearized about the a priori state to form the linearized measurements

$$\begin{aligned}z_\rho^{i0}(k) &= \nabla\Delta\rho_{AB}^{i0}(k) - \nabla\Delta\bar{r}_{AB}^{i0}(k) \\ &= (\hat{r}_B^0(k) - \hat{r}_B^i(k))^T \delta\vec{r}(k) + \nabla\Delta\nu_\rho \\ z_\phi^{i0}(k) &= \lambda_{L1}\nabla\Delta\phi_{AB}^{i0}(k) - \nabla\Delta\bar{r}_{AB}^{i0}(k) \\ &= (\hat{r}_B^0(k) - \hat{r}_B^i(k))^T \delta\vec{r}(k) + \lambda_{L1}N_{AB}^{i0} \\ &\quad + \nabla\Delta\nu_\phi\end{aligned}\quad (7)$$

where $z_\rho^{i0}(k)$ and $z_\phi^{i0}(k)$ are the linearized double-differenced pseudorange and carrier phase measurements, $\nabla\Delta\bar{r}_{AB}^{i0}(k)$ is the a priori expected double-differenced range given by Eq. 6 using $\vec{r}_B(k)$ in place of $\vec{r}_B(k)$, and $\hat{r}_B^i(k)$ is the unit vector pointing from the mobile receiver to satellite i .

A.2 Matrix Equation Formation

If both receivers are simultaneously tracking the same $M+1$ satellites, then M linearized double-differenced

measurements are obtained at each time step of the form given in Eq. 7. Gathering these M equations into matrix form yields

$$\begin{bmatrix} \tilde{z}_\rho(k) \\ \tilde{z}_\phi(k) \end{bmatrix} = \begin{bmatrix} H_{\rho x}(k) & 0 \\ H_{\phi x}(k) & H_{\phi N} \end{bmatrix} \begin{bmatrix} \delta\vec{x}(k) \\ \vec{N} \end{bmatrix} + \begin{bmatrix} \nabla\Delta\tilde{v}_\rho \\ \nabla\Delta\tilde{v}_\phi \end{bmatrix} \quad (8)$$

where $\delta\vec{x}(k)$ is the a posteriori correction to the real-valued component of the state and

$$\begin{aligned} H_{\rho x}(k) &= H_{\phi x}(k) \\ &= \begin{bmatrix} (\hat{r}_B^0(k) - \hat{r}_B^1(k))^T & 0 & 0 \\ \vdots & \vdots & \vdots \\ (\hat{r}_B^0(k) - \hat{r}_B^M(k))^T & 0 & 0 \end{bmatrix} \end{aligned} \quad (9)$$

$$H_{\phi N} = \lambda_{L1} \begin{bmatrix} 1 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ \vdots & & \ddots & \vdots \\ 0 & \dots & & 1 \end{bmatrix} \quad (10)$$

Assuming that pseudorange and carrier phase measurement noise for each satellite and receiver is independent and identically distributed, the measurement covariance matrices are easily derived as

$$\begin{aligned} R_\rho &= \sigma_\rho^2 \begin{bmatrix} 4 & 2 & \dots & 2 \\ 2 & 4 & \dots & 2 \\ \vdots & & \ddots & \vdots \\ 2 & \dots & & 4 \end{bmatrix} = \tilde{R}_\rho^{-T} \tilde{R}_\rho^{-1} \\ R_\phi &= \sigma_\phi^2 \begin{bmatrix} 4 & 2 & \dots & 2 \\ 2 & 4 & \dots & 2 \\ \vdots & & \ddots & \vdots \\ 2 & \dots & & 4 \end{bmatrix} = \tilde{R}_\phi^{-T} \tilde{R}_\phi^{-1} \end{aligned} \quad (11)$$

where σ_ρ and σ_ϕ are the standard deviation of the pseudorange and carrier phase measurements respectively and \tilde{R}_ρ and \tilde{R}_ϕ are the cholesky factorizations of the covariance matrices.

The state and state covariance are encoded as a set of equations in a square-root information filter. These equations are given as

$$\begin{bmatrix} \bar{\tilde{z}}_x(k) \\ \hat{\tilde{z}}_N(k-1) \end{bmatrix} = \begin{bmatrix} \bar{R}_{xx}(k) & \bar{R}_{xN}(k) \\ 0 & \hat{R}_{NN}(k-1) \end{bmatrix} \times \begin{bmatrix} \bar{\tilde{x}}(k) \\ \vec{N} \end{bmatrix} + \bar{w}_X \quad (12)$$

where \bar{w}_X is distributed as a multi-dimensional standard normal distribution, the state is determined from these equations through the solution process described in Sec. V-A.3, and the state covariance is given by

$$\begin{aligned} \bar{P}(k) &= \begin{bmatrix} \bar{R}_{xx}(k) & \bar{R}_{xN}(k) \\ 0 & \hat{R}_{NN}(k-1) \end{bmatrix}^{-T} \\ &\times \begin{bmatrix} \bar{R}_{xx}(k) & \bar{R}_{xN}(k) \\ 0 & \hat{R}_{NN}(k-1) \end{bmatrix}^{-1} \end{aligned} \quad (13)$$

The square-root information equations from Eq. 12 are appended to the measurement equations from Eq. 8 to incorporate the a priori information about the state into the solution process. The augmented update equations are normalized by the noise covariance matrix and given as

$$\begin{aligned} \begin{bmatrix} \tilde{\tilde{z}}_\rho(k) \\ \tilde{\tilde{z}}_\phi(k) \\ \delta\tilde{\tilde{z}}_x(k) \\ \tilde{\tilde{z}}_N(k) \end{bmatrix} &= \begin{bmatrix} \tilde{R}_\rho \tilde{z}_\rho(k) \\ \tilde{R}_\phi \tilde{z}_\phi(k) \\ \bar{\tilde{z}}_x(k) - \bar{R}_{xx}(k) \bar{\tilde{x}}(k) \\ \hat{\tilde{z}}_N(k-1) \end{bmatrix} \\ &= \begin{bmatrix} \tilde{H}_{\rho x}(k) & 0 \\ \tilde{H}_{\phi x}(k) & \tilde{H}_{\phi N} \\ \bar{R}_{xx}(k) & \bar{R}_{xN}(k) \\ 0 & \bar{R}_{NN}(k) \end{bmatrix} \begin{bmatrix} \delta\vec{x} \\ \vec{N} \end{bmatrix} \\ &+ \bar{w} \\ &= \begin{bmatrix} \tilde{R}_\rho H_{\rho x}(k) & 0 \\ \tilde{R}_\phi H_{\phi x}(k) & \tilde{R}_\phi H_{\phi N} \\ \bar{R}_{xx}(k) & \bar{R}_{xN}(k) \\ 0 & \hat{R}_{NN}(k-1) \end{bmatrix} \\ &\times \begin{bmatrix} \delta\vec{x} \\ \vec{N} \end{bmatrix} + \bar{w} \end{aligned} \quad (14)$$

where \bar{w} is distributed as a multi-dimensional standard normal distribution. The matrix coefficient of the state vector is then factorized into an orthonormal matrix and an upper-right triangular matrix via QR factorization. This results in

$$\begin{bmatrix} \tilde{H}_{\rho x}(k) & 0 \\ \tilde{H}_{\phi x}(k) & \tilde{H}_{\phi N} \\ \tilde{R}_{xx}(k) & \tilde{R}_{xN}(k) \\ 0 & \tilde{R}_{NN}(k) \end{bmatrix} = \hat{Q}(k) \begin{bmatrix} \hat{R}_{xx}(k) & \hat{R}_{xN}(k) \\ 0 & \hat{R}_{NN}(k) \\ 0 & 0 \end{bmatrix} \quad (15)$$

The measurement equations are transformed using the orthonormal matrix from the factorization such that the measurement equations become

$$\begin{aligned} \begin{bmatrix} \delta \hat{z}_x(k) \\ \hat{z}_N(k) \\ \vec{z}_r(k) \end{bmatrix} &= \hat{Q}^T(k) \begin{bmatrix} \tilde{z}_\rho(k) \\ \tilde{z}_\phi(k) \\ \delta \tilde{z}_x(k) \\ \tilde{z}_N(k) \end{bmatrix} \\ &= \begin{bmatrix} \hat{R}_{xx}(k) & \hat{R}_{xN}(k) \\ 0 & \hat{R}_{NN}(k) \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \delta \vec{x} \\ \vec{N} \end{bmatrix} \\ &+ \hat{w} \end{aligned} \quad (16)$$

A.3 Solution Procedure

To determine the optimal set of integer ambiguities, an integer least-squares algorithm is applied. Integer least-squares algorithms solve the problem of minimizing $\|\hat{z}_N(k) - \hat{R}_{NN}(k)\vec{N}\|$ under the constraint that the solution, \vec{N}_{opt} , be a set of integers. Two popular algorithms for solving integer least-squares problems are the LLL algorithm, described in Ref. [8], and the LAMBDA method, described in Ref. [9]. The LLL algorithm was chosen for two reasons. First, a tight, easy-to-implement lower bound on the probability that the integers were fixed correctly was readily available, as presented in Ref. [8]. Second, a software package to perform the LLL algorithm, called MILES [10], was readily available.

Once the optimal integer ambiguity set has been determined using MILES, the a posteriori correction to the real-valued component of the state is determined as

$$\delta \vec{x}(k) = \hat{R}_{xx}^{-1}(k) \left(\delta \hat{z}_x(k) - \hat{R}_{xN}(k) \vec{N}_{opt} \right) \quad (17)$$

The a posteriori state covariance matrix can also be computed as

$$\begin{aligned} P(k) &= \begin{bmatrix} \hat{R}_{xx}(k) & \hat{R}_{xN}(k) \\ 0 & \hat{R}_{NN}(k) \end{bmatrix}^{-T} \\ &\times \begin{bmatrix} \hat{R}_{xx}(k) & \hat{R}_{xN}(k) \\ 0 & \hat{R}_{NN}(k) \end{bmatrix}^{-1} \end{aligned} \quad (18)$$

B. INS Propagation Step

The INS propagation step is responsible for propagating the state forward in time using the IMU's accelerometer measurements and attitude estimate. As mentioned previously, the attitude estimate provided by the IMU is assumed to be perfect for the purposes of this prototype system. This assumption was made to simplify development of the EKF and does not significantly effect the performance of the system due to the high accuracy of the IMU's attitude estimate.

Since the IMU relies on a magnetometer to determine heading, the IMU must be given the magnetic declination. This is accomplished by initializing the IMU with a rough estimate of position, which it uses in a magnetic field model to determine declination. If the system stays within a few kilometers of the initialization point, then this value will remain reasonably accurate due to the slow spacial variation of the magnetic field.

B.1 Dynamics and Measurement Models

The propagation step is essentially a double integrator of the accelerometer measurements after removing specific forces or accelerations the IMU senses other than motion in the ECEF frame. These other sensed components of the accelerometer measurements are the specific normal force and the Coriolis acceleration due to rotation of the Earth. The accelerometer may also have small biases in its measurements due to electrical biases in the circuitry. These biases will be modeled as a random walk process. This model for the accelerometer errors is similar to that presented in Ref. [11]. This yields the following continuous-time dynamics and measurement model

$$\begin{aligned}
\frac{d}{dt}\vec{r} &= \dot{\vec{r}} \\
\frac{d}{dt}\dot{\vec{r}} &= R_{ECEF}^{IMU}(\vec{f} - \vec{b}) + R_{ECEF}^{ENU} \begin{bmatrix} 0 \\ 0 \\ -g \end{bmatrix} \\
&\quad - 2\vec{\omega}_E \times \dot{\vec{r}} + \vec{v}_a \\
\frac{d}{dt}\vec{b} &= \vec{v}_b \\
\frac{d}{dt}\vec{N} &= 0
\end{aligned} \tag{19}$$

where \vec{f} is the measured specific force vector (i.e. accelerometer measurement), R_{ECEF}^{ENU} is the rotation matrix from the east-north-up (ENU) frame to the ECEF frame that is derived from the a priori position estimate of the IMU, g is the gravitational acceleration of Earth at the location of the IMU, $\vec{\omega}_E$ is the angular velocity vector of the Earth, and \vec{v}_a and \vec{v}_b are zero-mean, white noise processes with standard deviations of σ_a and σ_b respectively.

Equation 19 is discretized using an Euler integration approximation to determine the state transition matrix, which assumes a small time step between updates. Since the IMU reports measurements every 0.01 s, this approximation is reasonable. The integer ambiguity vector does not change with time (i.e. $\vec{N}(k+1) = \vec{N}(k)$), so its propagation equation is trivial and thus eliminated from the propagation step. The discrete-time propagation model is given in matrix form as

$$\begin{aligned}
\vec{x}(k+1) &= \begin{bmatrix} \vec{r}(k+1) \\ \dot{\vec{r}}(k+1) \\ \vec{b}(k+1) \end{bmatrix} \\
&= \begin{bmatrix} I & I\Delta t & 0 \\ 0 & I - 2[\vec{\omega}_E \times] \Delta t & -R_{ECEF}^{IMU} \Delta t \\ 0 & 0 & I \end{bmatrix} \\
&\quad \times \begin{bmatrix} \vec{r}(k) \\ \dot{\vec{r}}(k) \\ \vec{b}(k) \end{bmatrix} + \begin{bmatrix} 0 \\ R_{ECEF}^{IMU} \vec{f} + \vec{g}_{ECEF} \\ 0 \end{bmatrix} \Delta t + \vec{v}' \\
&= F(k)\vec{x}(k) + \vec{u}(k) + \vec{v}'
\end{aligned} \tag{20}$$

where $[\vec{\omega}_E \times]$ is the cross-product-equivalent matrix for the angular velocity of the Earth, $\vec{g}_{ECEF} = \left(\begin{bmatrix} 0 & 0 & -g \end{bmatrix} (R_{ECEF}^{ENU})^T \right)^T$, and \vec{v}' is the process

noise vector. It can be shown that the process noise covariance matrix is

$$Q = \begin{bmatrix} 0 & 0 & 0 \\ \frac{1}{3}\sigma_a^2 \Delta t^3 I & \frac{1}{2}\sigma_a^2 \Delta t^2 I & 0 \\ \frac{1}{2}\sigma_a^2 \Delta t^2 I & I(\sigma_a^2 \Delta t + \frac{1}{2}\sigma_b^2 \Delta t^2) & -\frac{1}{2}R_{ECEF}^{IMU} \sigma_b^2 \Delta t^2 \\ 0 & -\frac{1}{2}R_{ECEF}^{IMU} \sigma_b^2 \Delta t^2 & \sigma_b^2 \Delta t I \end{bmatrix} \tag{21}$$

B.2 Square-Root Information Propagation

For use in the EKF, the discrete-time propagation model in Eq. 20 is converted into square-root information form. Performing the propagation in this form aids in preserving the information relating the real and integer portions of the state through the square-root information equations given in Eq. 12. Starting with the square-root information equations, the discrete-time propagation model is used to substitute the state at time $k+1$ for the state at time k . The resulting equations, after augmenting with the square-root equation for the process noise, are

$$\begin{aligned}
&\begin{bmatrix} 0 \\ \vec{z}_x(k) + R_{xx}(k)F^{-1}(k)\vec{u}(k) \\ \vec{z}_N(k) \end{bmatrix} = \\
&\begin{bmatrix} R_{vv} & 0 & 0 \\ -R_{xx}(k)F^{-1}(k) & R_{xx}(k)F^{-1}(k) & R_{xN}(k) \\ 0 & 0 & R_{NN}(k) \end{bmatrix} \\
&\quad \times \begin{bmatrix} \vec{v}' \\ \vec{x}(k+1) \\ \vec{N} \end{bmatrix} + \vec{w}_{v,X}
\end{aligned} \tag{22}$$

where R_{vv}^T is the Cholesky factorization of the process noise covariance matrix and $\vec{w}_{v,X}$ is distributed as a multi-dimensional standard normal distribution. As with the CDGPS measurement update, the matrix coefficient of the state vector is then factorized into an orthonormal matrix and an upper-right triangular matrix via QR factorization. This results in

$$\begin{aligned}
& \begin{bmatrix} R_{vv} & 0 & 0 \\ -R_{xx}(k)F^{-1}(k) & R_{xx}(k)F^{-1}(k) & R_{xN}(k) \\ 0 & 0 & R_{NN}(k) \end{bmatrix} \\
& = \bar{Q}(k) \begin{bmatrix} \bar{R}_{vv}(k+1) & \bar{R}_{vx}(k+1) & \bar{R}_{vN}(k+1) \\ 0 & \bar{R}_{xx}(k+1) & \bar{R}_{xN}(k+1) \\ 0 & 0 & \bar{R}_{NN}(k+1) \end{bmatrix} \quad (23)
\end{aligned}$$

The propagation equations are transformed using the orthonormal matrix from the factorization such that they become

$$\begin{aligned}
& \begin{bmatrix} \bar{z}_v(k+1) \\ \bar{z}_x(k+1) \\ \bar{z}_N(k+1) \end{bmatrix} = \\
& \bar{Q}^T(k) \begin{bmatrix} 0 \\ \bar{z}_x(k) + R_{xx}(k)F^{-1}(k)\bar{u}(k) \\ \bar{z}_N(k) \end{bmatrix} \\
& = \begin{bmatrix} \bar{R}_{vv}(k+1) & \bar{R}_{vx}(k+1) & \bar{R}_{vN}(k+1) \\ 0 & \bar{R}_{xx}(k+1) & \bar{R}_{xN}(k+1) \\ 0 & 0 & \bar{R}_{NN}(k+1) \end{bmatrix} \quad (24) \\
& \times \begin{bmatrix} \bar{v}' \\ \bar{x}(k+1) \\ \bar{N} \end{bmatrix} + \bar{w}_{v,X}
\end{aligned}$$

The a priori state and state covariance can be determined in the same way as described in Sec. V-A.3. The propagated square-root information equations are passed on to the next operation, which could be a CDGPS measurement update or another IMU propagation step depending on what measurements are next. Under the current configuration of the PAR system, GPS observables are reported every 0.2 s and IMU measurements and attitude are reported every 0.01 s. This leads to 20 sequential propagation steps being performed between every measurement update.

C. Augmented Reality Overlay

The augmented reality overlay is responsible for taking the precise position and attitude of the webcam provided by the EKF, capturing the webcam's view in a virtual world referenced to real world coordinates, and overlaying the view of the virtual world on top of the webcam footage, which constitutes the user's

view of the real world. This combination of real world video footage and virtual images represented in their assigned real-world location and attitude comprises the augmented reality experience. The virtual world could contain any objects or text for any purpose, which makes the PAR system application ambivilous.

For the prototype PAR system, MATLAB's SIMULINK 3D Animation Toolbox was employed to create and visualize the virtual world and perform the overlay of the view of the virtual world on top of the real world footage. The overlay was performed by adding an appropriately sized, positioned, and oriented box in the virtual world that acted as a "projection screen" for the webcam footage, which was applied to the box as a texture. While this method for performing the image overlay has its drawbacks when it comes to adding lighting and shadowing effects to the virtual world, the method worked well for the prototype PAR system, which did not incorporate these effects. An example of this overlay is shown in Fig. 8. The image on the left side of the figure shows the view of a virtual piston before the overlay, and the image on the right side shows the result of the overlay.

VI. TEST RESULTS

The prototype augmented reality system was tested under two different scenarios. The first test was a static test that was meant to verify that the CDGPS algorithm was performing as expected. The second test was a dynamic test that was meant to test the full functionality of the system. For ease of testing, all processing was done after-the-fact by operating on recordings of the IMU data, GPS observables, and webcam video.

A. Static Test

Before the start of the static test, the baseline distance between the base station antenna and the secondary antenna used for the static test was measured using a tape measure, as seen in Fig. 9. This measurement provided the truth data for the static test to verify that the CDGPS algorithm was converging correctly. The baseline distance measured was 21.155 m.

30 minutes of data was collected and processed for the static test. The base station and secondary receiver

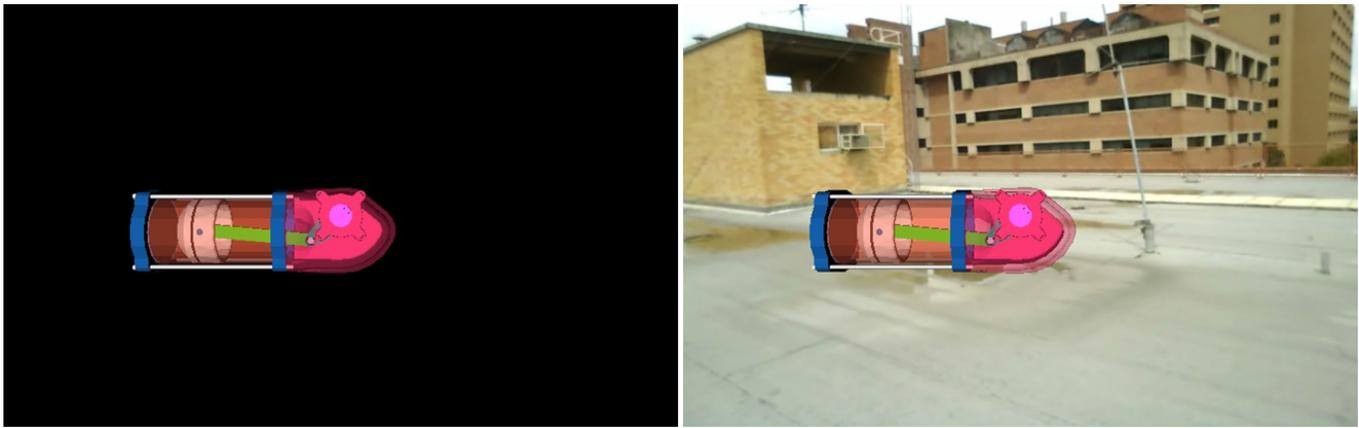


Fig. 8. An example of the overlay of the virtual world and real world footage. The image on the left side shows the view of a virtual piston before the overlay. The image on the right side shows the result of the overlay.

tracked between 8 and 10 common satellites at various points during the 30 minute interval, as shown in Fig. 10. The dropping and adding of satellites throughout the dataset demonstrates the CDGPS algorithm’s capability to operate through changes in the satellites tracked without any issues. This was enabled by special modifications that the CDGPS algorithm makes to the square-root information equations whenever a satellite is lost or acquired.

Figure 11 shows a plot of the lower bound on the probability that the integer ambiguities have been fixed correctly during the first portion of the static test. It can be assumed with high certainty that the integers have converged correctly once this probability has exceeded 0.999, which corresponds to less than a 0.1% probability of error. Using this metric, the correct integer ambiguities were determined within the first 11.2 s of the dataset.

Figure 12 shows the position estimate of the secondary antenna in the ENU coordinate system centered on the base station antenna from the time that the probability metric was exceeded to the end of the test. Almost all of the points in the plot fit within a centimeter range in both East and North directions. The baseline distance was calculated from these points and is shown in Fig. 13. The mean baseline distance was 21.1587 m, which is within a few millimeters of the value measured with a tape measure. This demonstrates that the CDGPS algorithm is performing as expected.

B. Dynamic Test

The data for the dynamic test was collected over a few minutes. For the first 2–3 minutes, the prototype system was stationary on a ledge to allow the integer ambiguities to converge. Then, the prototype system was walked toward and around a virtual object, which would later be visualized in post-processing. Finally, the system was returned to its starting position on the ledge.

Figure 14 shows the satellites tracked by both the base station and mobile receivers during the entire dataset. In this case, the same 7 satellites were tracked throughout the dataset due to the short duration of the test. Figure 15 shows the lower bound probability on correct convergence of the integer ambiguities. In this case, the probability of correct resolution exceeded the 0.999 metric after 31.8 s.

Figure 16 shows the position estimate of the mobile antenna in the ENU coordinate system centered on the base station antenna from the time that the probability metric was exceeded to the end of the test. The position estimates are good enough that the wobbles in the trajectory as each step was taken can be clearly seen in the plot. It can also be seen that the position estimate did not drift during the test, since the position estimate at the start and end of the test are nearly identical. The IMU’s orientation estimate throughout the test is shown in Fig. 17 as a roll-pitch-yaw Euler angle sequence. Due to the orientation of the IMU in the prototype system, the roll-pitch-yaw sequence if one were to hold the system level to the ground and point the camera to the North would be a 90° roll, 0° pitch, and -90° yaw.

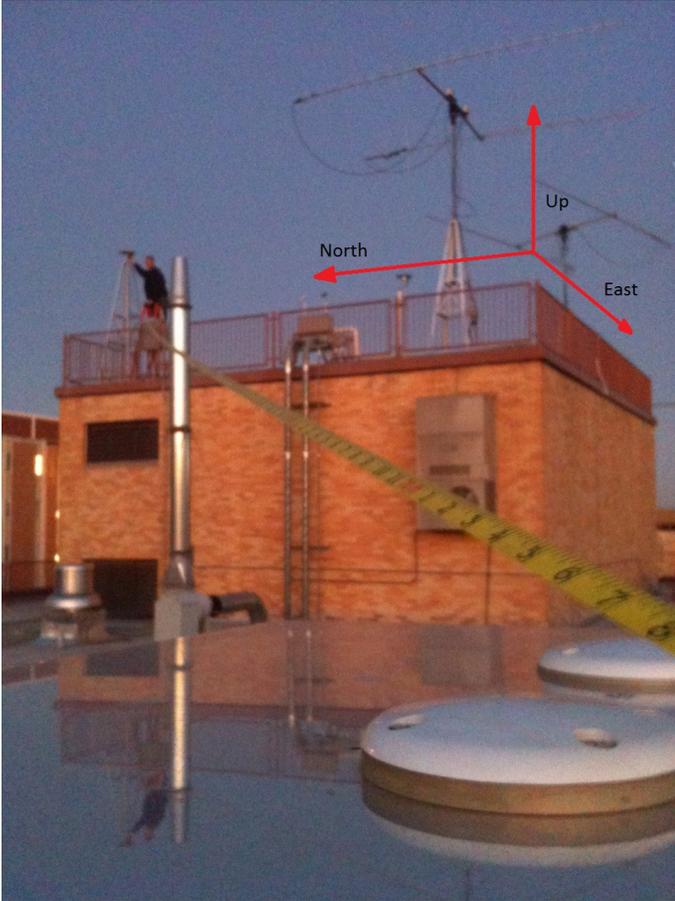


Fig. 9. Picture of measuring the baseline distance for the static test using a tape measure.

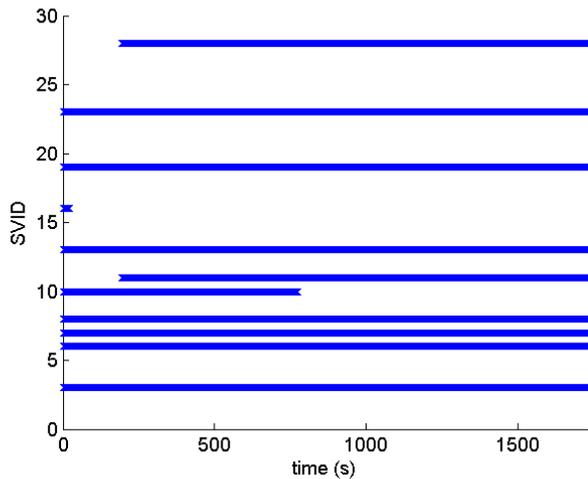


Fig. 10. The SVIDs of the GPS satellites tracked during the static test as a function of time.

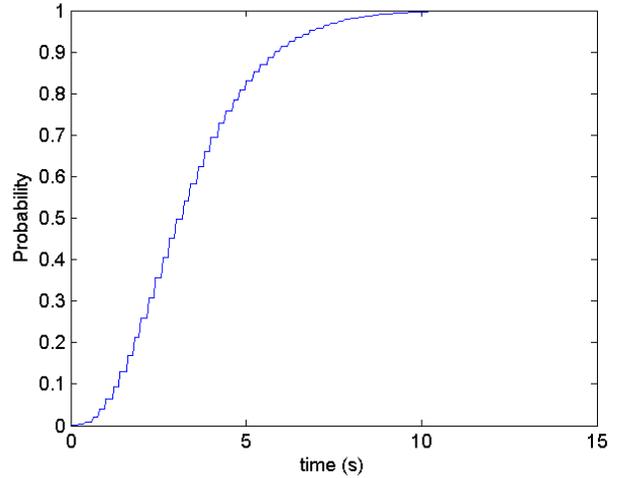


Fig. 11. The lower bound on the probability that the integer ambiguities have been fixed correctly as a function of time for the static test.

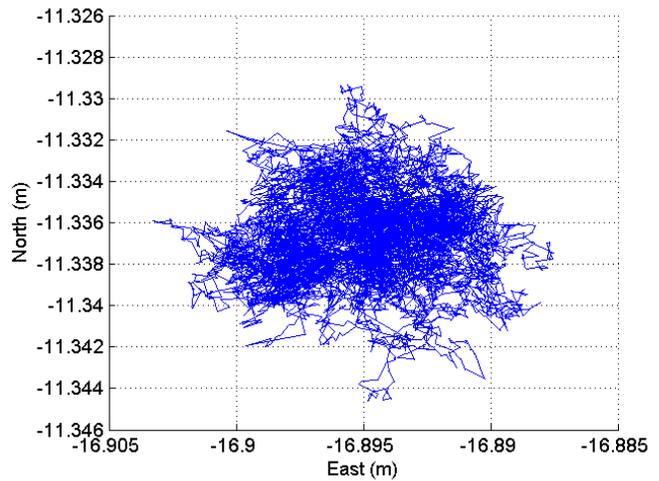


Fig. 12. A plot of the East and North position relative to the base station during the static test.

A virtual world was designed to demonstrate the augmented reality overlay portion of the system. This virtual world contained a set of blue marble columns that the user passed through twice during the test and a model of a piston that the user walks up to, around, and then backs away from. Figure 18 shows a sequence of frames taken from the augmented reality video that was produced using the described virtual world, the webcam video, and the output position and attitude of the webcam from the EKF. The rendering of the virtual world was incredibly stable relative to the real world and contained little so-called registration error.

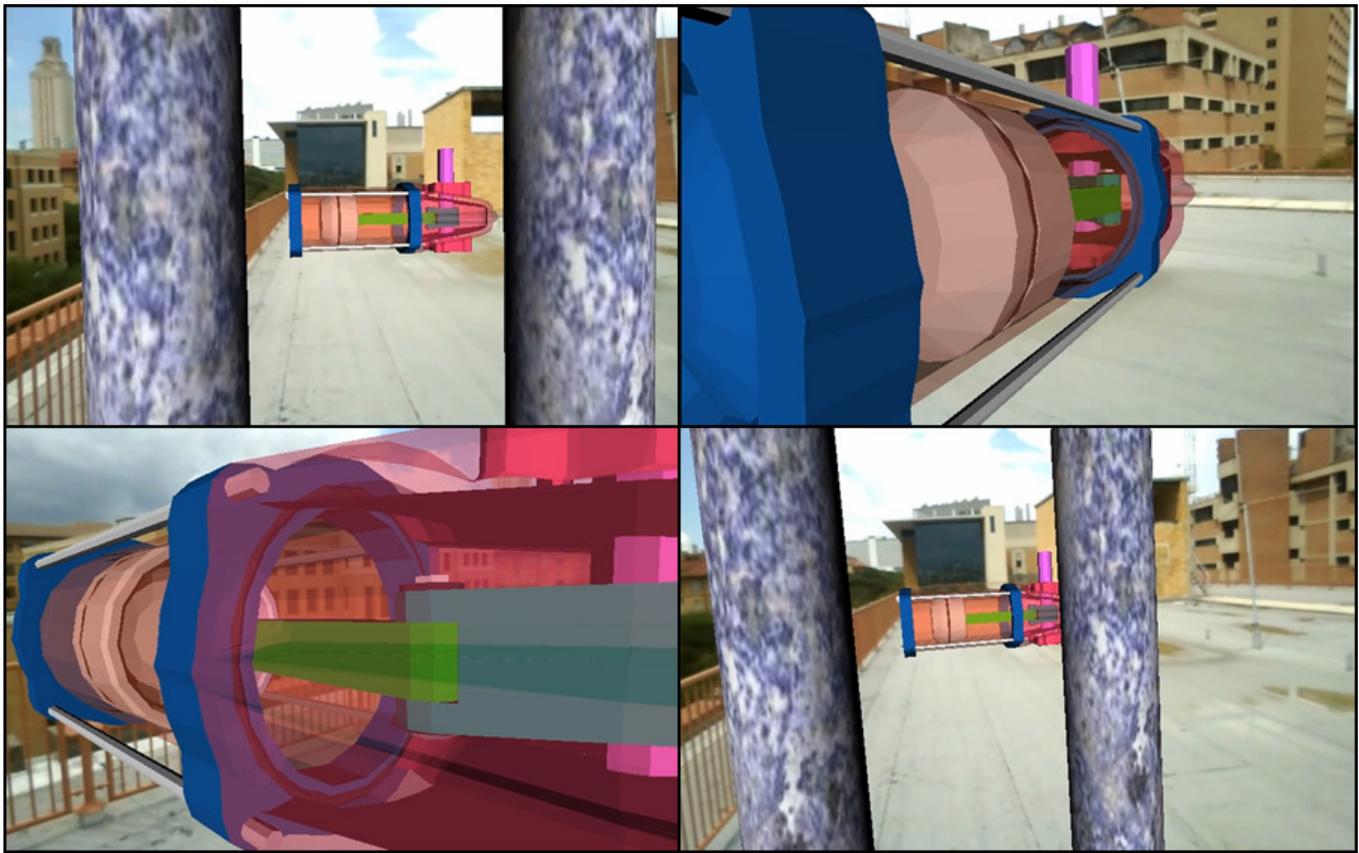


Fig. 18. A sequence of frames taken from the augmented reality video produced using the results from the dynamic test. The first frame (upper left) was taken from shortly after the system first started moving. The second frame (upper right) was taken from the moment the system reaches the southern-most point of its walk around the piston. The third frame (lower left) was taken from the moment the system reaches the northern-most point of its walk around the piston. The fourth frame (lower right) was taken from near the end of the dataset.

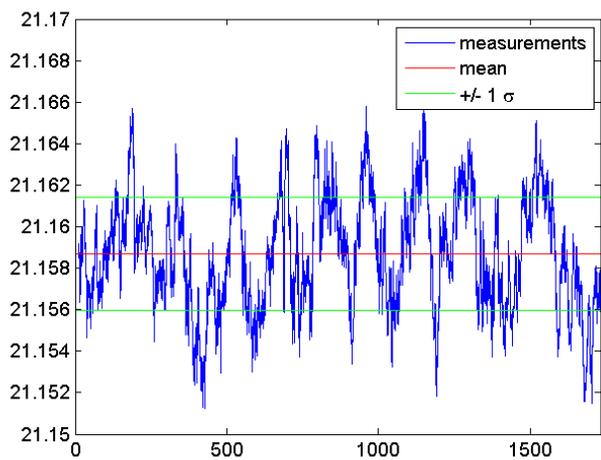


Fig. 13. A plot of the measured baseline distance over the 30 minute static test. The red line is the mean distance and the green lines mark the $\pm 1\sigma$ bounds.

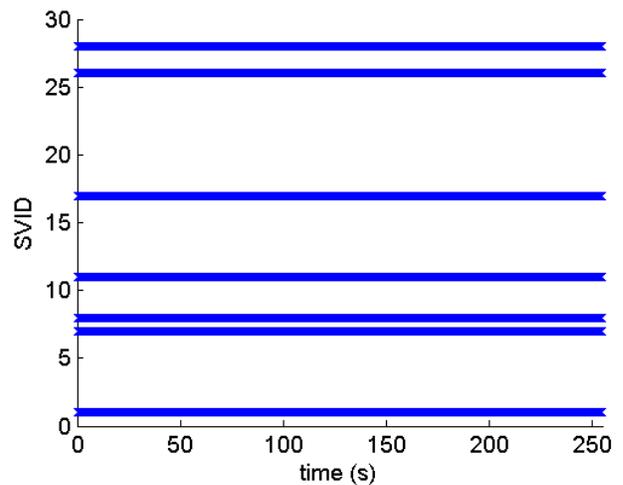


Fig. 14. The SVIDs of the GPS satellites tracked during the dynamic test as a function of time.

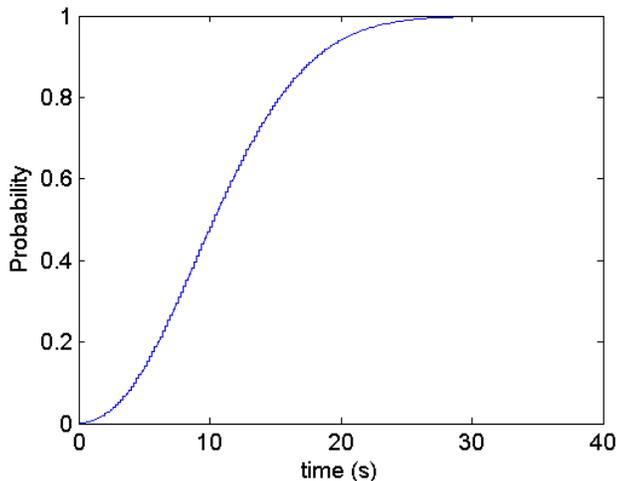


Fig. 15. The lower bound on the probability that the integer ambiguities have been fixed correctly as a function of time for the dynamic test.

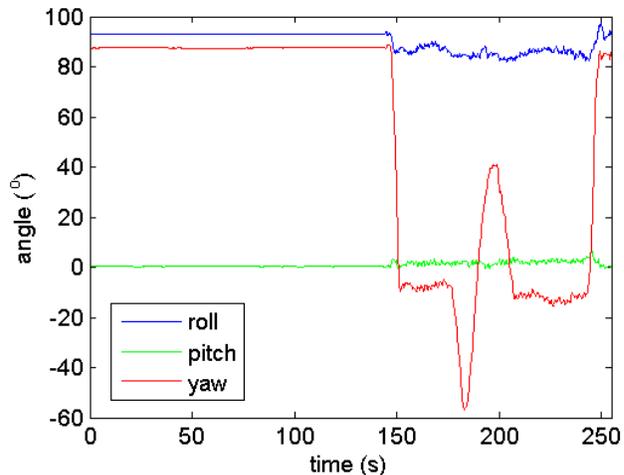


Fig. 17. A plot of the Euler angles reported by the IMU during the dynamic test. The IMU coordinate system is aligned such that the x-axis points to the right and the y-axis points up as seen by a user holding the system.

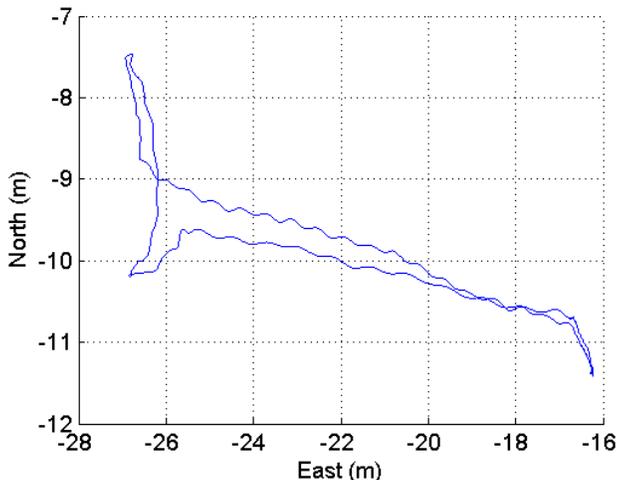


Fig. 16. A plot of the East and North position relative to the base station during the dynamic test.

VII. CHALLENGES / FUTURE WORK

This section provides a brief description of some of the technical difficulties associated with translating this demonstration and prototype PAR system into a viable commercial product.

A. Carrier Phase Cycle Slips

The integer ambiguities solved for by the CDGPS algorithm are only good so long as the GPS receiver's PLL is able to track the signal without slipping any cycles in the carrier phase. These cycle slips can occur due to receiver dynamics, low carrier-to-noise ratio,

Ionospheric changes, and multipath. If the receiver slips a cycle, then the CDGPS algorithm would no longer report the correct position unless this cycle slip is detected and mitigated. Properly designed receivers will experience few cycle slips, but cycle slips cannot be completely eliminated. Research has been done on this topic in Ref. [7], but further research is warranted for this particular application of cycle slip detection and mitigation.

B. Carrier Phase Wind-up

Carrier phase wind-up is a phenomenon caused by the circular polarization of the signal. If the GPS antenna is rotated, then the received carrier phase of each signal will increase or decrease, depending on the direction of rotation and polarization of the signal, by an amount proportional to the angle rotated. This will lead to errors in the carrier phase measurements that will drift in time in the case of human motion. However, the IMU keeps track of the orientation of the PAR system, so, given a model for carrier phase wind-up based on antenna orientation, the PAR system should be able to eliminate the effects of carrier phase wind-up. One such model for carrier phase wind-up is presented in Ref. [12].

C. GPS Reception

The problem of GPS signal reception encompasses three different problems, antenna design, self-interference, and indoor and deep-urban navigation. Incorporating a GPS antenna into a small device with good phase properties and reception gain pattern is a difficult challenge that poses a major obstacle to integrating CDGPS capability into a device like a smartphone. One possible solution to this problem is to accept that the system would have to be either a stand-alone device or a slightly larger attachable accessory. Self-interference refers to the problem of isolating the GPS antenna from any potential interference sources within the PAR system itself. Oftentimes clocks, processors, and displays transmit noise in the GPS band that, in close proximity, can cause issues with GPS reception.

Indoor and deep-urban navigation also pose significant challenges to satellite positioning due to the necessarily weak signals [13]. The surveying community deals with deep-urban environments by using multiple satellite constellations and frequencies, but this functionality comes with a significantly elevated price tag. Due to this problem, any PAR system would likely only be able to function in areas with a clear view of the sky. This constraint, however, does not prevent use of a PAR system in many potential applications.

VIII. CONCLUSIONS

Precise augmented reality (PAR) will open the door for a whole host of augmented reality applications by enabling alignment of the virtual and real worlds at short distances (< 100 m) without requiring visual cues. Most current augmented reality systems are only capable of positioning accuracies in the 3-to-10 meter range, while a PAR system is capable of centimeter and even sub-centimeter level accuracies. Only the absolute positioning accuracy of a PAR system is capable of reducing the alignment or “registration” errors to the level required for small relative distances between the user and virtual objects without visual cues.

A prototype PAR system was designed, built, and tested in both static and dynamic scenarios. The static test demonstrated that the CDGPS algorithm was providing sub-centimeter level accurate position-

ing, as expected. The dynamic test demonstrated the accuracy of the system in both position and orientation as it was walked around a virtual object. The rendering of the virtual object was remarkably stable in position and orientation relative to the real world camera feed.

While the results of these tests are promising, there are still several issues that need to be resolved to make the algorithm’s performance robust. These issues include cycle slip detection and mitigation and carrier phase wind-up. GPS reception is also an obstacle which inhibits performance of the system. This issue manifests itself in several ways including antenna design, self-interference, and the difficulty of indoor and deep-urban navigation. These issues are not likely to be insurmountable, with the possible exception of indoor and deep-urban navigation. However, there are still many potential applications where sky visibility is not an issue that would benefit from a PAR system.

References

- [1] Sportsvision, “1st and ten line system,” <http://www.sportvision.com/foot-1st-and-ten-line-system.html>, 2012, [Online; accessed 27-September-2012].
- [2] Lego, “Lego augmented reality kiosks heading to shops worldwide,” 2010.
- [3] V. Technology, “Start walk,” <http://vitotechnology.com/star-walk.html>, 2012, [Online; accessed 28-September-2012].
- [4] G. Roberts, A. Evans, A. Dodson, B. Denby, S. Cooper, R. Hollands *et al.*, “The use of augmented reality, gps and ins for subsurface data visualization,” in *FIG XXII International Congress*, 2002, pp. 1–12.
- [5] S. Feiner, B. MacIntyre, T. Höllerer, and A. Webster, “A touring machine: Prototyping 3d mobile augmented reality systems for exploring the urban environment,” *Personal and Ubiquitous Computing*, vol. 1, no. 4, pp. 208–217, 1997.
- [6] B. O’Hanlon, M. Psiaki, S. Powell, J. Bhatti, T. E. Humphreys, G. Crowley, and G. Bust, “CASES: A smart, compact GPS software receiver for space weather monitoring,” in *Proceedings of the ION GNSS Meeting*. Portland, Oregon: Institute of Navigation, 2011.
- [7] S. Mohiuddin and M. L. Psiaki, “High-altitude satellite relative navigation using carrier-phase differential global positioning system techniques,” *Journal of Guidance, Control, and Dynamics*, vol. 30, no. 5, pp. 1628–1639, Sept.–Oct. 2007.
- [8] A. Hassibi and S. Boyd, “Integer parameter estimation in linear models with applications to GPS,” *IEEE Transactions on Signal Processing*, vol. 46, no. 11, 1998.
- [9] M. Psiaki and S. Mohiuddin, “Global positioning system integer ambiguity resolution using factorized least-squares techniques,” *Journal of Guidance, Control, and Dynamics*, vol. 30, no. 2, pp. 346–356, March–April 2007.
- [10] X. W. Chang, X. Xie, and T. Zhou, *MILES: MATLAB package for solving Mixed Integer LEast Squares prob-*

lems, 2nd ed., <http://www.cs.mcgill.ca/chang/software.php>, Oct. 2011.

- [11] W. S. Flenniken IV, J. H. Wall, and D. M. Bevly, "Characterization of various imu error sources and the effect on navigation performance," in *Proceedings of the ION ITM*. Long Beach, California: Institute of Navigation, 2005.
- [12] J. Wu, S. Wu, G. Hajj, W. Bertiguer, and S. Lichten, "Effects of antenna orientation on gps carrier phase measurements," *Manuscripta Geodaetica*, pp. pp. 91–98, 1993.
- [13] G. Seco-Granados, J. Lopez-Salcedo, D. Jimenez-Banos, and G. Lopez-Risueno, "Challenges in indoor global navigation satellite systems: Unveiling its core features in signal processing," *Signal Processing Magazine, IEEE*, vol. 29, no. 2, pp. 108–131, march 2012.